A Framework for Application-aware Networking by Delegating Traffic Management of SDNs

Hamidreza Almasi Kashi
Department of Computer Engineering
Sharif University of Technology
Tehran, Iran
Email: almasi@ce.sharif.edu

Hossein Ajorloo
School of Computer Science
Institute for Research in Fundamental Sciences (IPM)
Tehran, Iran
Email: ajorloo@ipm.ir

arXiv:1610.05062v1 [cs.NI] 17 Oct 2016

***Abstract*—Presence of a logically centralized controller in software-defined networks enables smart and fine-grained management of network traffic. Generally, traffic management includes measurement, analysis and control of traffic in order to improve resource utilization. This is done by inspecting corresponding performance requirements using metrics such as packet delay, jitter, loss rate and bandwidth utilization from global network view. There has been many works regarding traffic management of software-defined networks and how it could help to efficiently allocate resources. However, the vast majority of these solutions are bounded to indirect information retrieved within the border of ingress and egress switches. This means that the three stage loop of measurement, analysis and control is performed on switches in between this border while the traffic flowing in network originates from applications on end hosts. In this work, we present a framework for incorporating network applications into the task of traffic management using the concept of software-defined networking. We demonstrate how this could help applications to receive desired level of quality of service by implementing a prototype of an API for flow bandwidth reservation using OpenFlow and OVSDB protocols.**

I. Introduction

The term SDN was first introduced in an article which was describing the OpenFlow project at Stanford University [1]. OpenFlow stemmed from observing the fact that most new networking community ideas went untried and untested, so it was considered a means for researchers to let them test their experimental protocols in a production network [2]. It enabled programmability of network by setting rules in switches flow tables. Since then, the definition of SDN has expanded and OpenFlow is just one possible data plane or southbound API. Furthermore, SDN relies on control and data plane separation but does not reduce to just this, but in general, SDN is all about abstraction provided by forwarding devices with well-defined instruction sets, carefully designed network operating systems called controllers along with data and control plane separation by using appropriate southbound and northbound APIs [3].

The controller in software-defined networks maintains a global network view of traffic flows, link connectivity states, network resource allocation, etc. The control plane in these networks installs arbitrary forwarding and monitoring rules in TCAM entries of SDN switches to be matched on by data plane. This provides near-optimal, fine-grained decision making for traffic management tasks such as load balancing, fault tolerance, topology updates and traffic analysis compared to sub-optimal, coarse-grained decisions made in legacy networks.

Measuring and controlling QoS parameters like packet delay, jitter, loss rate and throughput is an inherent part of traffic management tasks and an important input to many decisions in computer networks. There has been several approaches towards maintaining these parameters in different types of networks. With more visibility into network resources, SDN introduced more efficient approaches [4]. However, these solutions are limited to knowledge of measurements performed on flow table entries within ingress and egress switches area. In other words, knowledge of controller about network traffic is confined to statistics of the rules it eventually had to install in order to forward the traffic towards destination while the traffic originates from end host applications.

In this paper, we argue that in some environments such as data centers, these network applications better know the nature and anticipated QoS of the traffic they are generating and thus might be able to help the controller to treat the traffic in a more efficient manner. When network application programmers are sending data in some parts of their code development process, they could ask the network controller through an API for special QoS treatment tailored to the nature of the data and the controller could investigate the possibility of providing the service and respond accordingly. By using this framework, some portions of traffic management is delegated to applications and the network becomes more application-aware. This is not straightly possible in legacy networks due to distributed control. We will discuss that DiffServ [5] does not provide fine-grained guarantees some applications might expect. However, the logically centralized controller in software-defined networks enables more direct monitoring and control on network resources while being a designated authority who the applications could talk to through an API.

The remainder of this paper is organized as follows: Section II discusses related work on traffic management. Section III presents the suggested framework and Section IV continues by presenting a use case implementation for guaranteed minimum bandwidth. Section V evaluates the prototype and Section VI concludes the paper.

II. Related Work

There are several research works in the field of quality of service maintenance and using application layer information both in legacy and software-defined networks. We tend to discuss more practical solutions here.

DiffServ [5] is a simple and scalable mechanism for classification and management of traffic to provide QoS in IP networks. This method can bring low-delay service for latency-sensitive traffic while using best-effort delivery for web and file transfer purposes. This class-based, coarse-grained mechanism, performs all traffic classification at DiffServ domains borders and conceals the complexity from network core and unlike IntServ [6] and RSVP [7], it does not require reservation and end-to-end signaling, but on the other hand, how each router deals with DSCP bits depends on its configuration and therefore it is hard to predict end-to-end behavior.

Hedera [8], is a dynamic and scalable solution for flow scheduling avoiding ECMP limitations. It has a global view of traffic demands and routing information. Hedera scheduler collects flow information from switches, computes collision free paths and instructs switches to conduct traffic along those paths in a control loop. This periodic polling of traffic information or packet sampling leads to high monitoring overheads which incurs significant switch resources consumption and long event detection times. Mahout [9] suggests to use an end host shim layer which marks the packets of elephant flows. The switches are configured to forward marked packets to the Mahout controller. This resolves switch monitoring overhead issue. MicroTE [10] notes that recent proposals for traffic engineering in ISPs operate at long time-scales and based on measurements [11], data center traffic is naturally busty and unpredictable at such long time-scales but a significant amount of traffic is predictable at short time-scales. It assigns servers for aggregating traffic statistics and sending summarized traffic matrix to the network controller which takes appropriate routing decisions for both predictable and non-predictable traffic. Yet, using either shim layer or top-of-rack aggregator may not deliver fine-grained control to individual network applications.

In [12], a QoS controller is presented which can create network slices and provision them dynamically to satisfy performance requirements across applications it has assigned their traffic to those slices. Network administrators specify high level slice specifications and controller reserves network resources accordingly. Their main difference to our approach is that we delegate traffic management of our software-defined network to its applications and network administrators specify high level policies and limits for those authorized applications rather than slice specifications and we do this by developing a northbound API rather than extending southbound APIs such as OpenFlow since there is potentially more work to do in this part.

III. Proposed Framework

As mentioned, participating network applications in traffic management task of a software-defined network could better satisfy their service performance requirements. Here, we present the suggested framework which enables the controller to be aware of how the developer expected the network to behave with the traffic coming from the application.

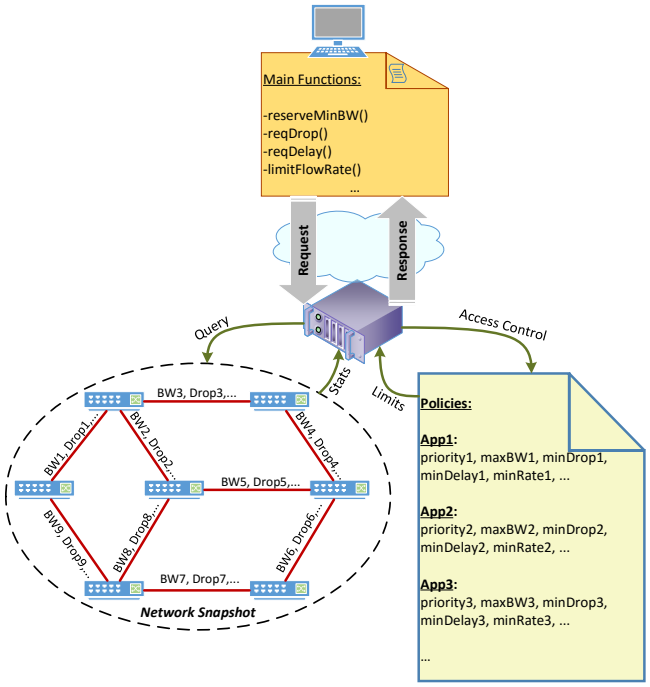

Fig. 1. Framework overview

A. Overview

Using this framework, network application developers use an API to connect to destination with a differentiated socket when they require special treatment from the network. When a function is called from this API, there will be a negotiation phase with the controller which is constantly updating a snapshot of network resources status. The controller would then accept, reject or revise the request and then implements the behavior as rules in case it conforms to policies provided by network administrator and required resources are available. Finally, in case of request acceptance, the controller reserves resources for that flow for some time and then the connection is made as usual. Otherwise it enforces a revision or rejects the request and the connection is made without any special treatment. Requests from applications are put in a queue and processed according to application priority or on a first-come, first-served basis in case of equal priorities. An overview of this framework with its abstract interactions is depicted in Fig. 1.

Four instances of API functions with their corresponding policies and network statistics data structures are shown in the figure. Their common arguments are `socketType`, `destIP` and `destPort`. Their specific argument and policy check expression is indicated in Table I. Method simplePaths(G, source, dest) operates over graph $G = (V, E)$ from a topology discovery module feeding network snapshot where V is the set

TABLE I
API SPECIFICATIONS

Function	Specific Argument	Policy Check Expression	Description
`reserveMinBW()`	`minExpectedBandwidth`	$(minExpectedBandwidth \leq maxBW) \wedge (\exists p \in simplePaths(G, source, dest) \mid \min_{\forall l \in L_p} (C_l - U_l) > minExpectedBandwidth)$	Guarantees minimum bandwidth a flow receives. L_p is the set of links forming path p. C_l and U_l are capacity and utilization of link l respectively.
`reqDrop()`	`maxExpectedDropRate`	$(maxExpectedDropRate \geq minDrop) \wedge (\exists p \in simplePaths(G, source, dest) \mid \max_{\forall i \in I_p} (D_i) \leq maxExpectedDropRate)$	Enforces an upper bound for the drop rate a flow might suffer. I_p is the set of network interfaces along path p. D_i is the drop rate of interface i.
`reqDelay()`	`maxExpectedDelay`	$(maxExpectedDelay \geq minDelay) \wedge (\exists p \in simplePaths(G, source, dest) \mid Delay(p) \leq maxExpectedDelay)$	Enforces an upper bound for the delay a flow might encounter. Delay(p) calculates average latency the flow packets encounter along path p.
`limitFlowRate()`	`maxFlowRate`	$(maxFlowRate \geq minRate) \wedge (\exists p \in simplePaths(G, source, dest) \mid U_l \underset{\forall l \in L_p}{<} C_l)$	Enforces traffic contract between tenants of flow source and destination. C_l and U_l are capacity and utilization of link l in path p respectively.

of OpenFlow switches and E is the set of links interconnecting them. Note that these functions are only some samples of how QoS enforcement tasks could be plugged into this framework and the generality of framework allows for other forms of functions and parameters. We next describe how the controller becomes aware of application intents for traffic.

B. Application-Controller Interaction

When an application decides to send some traffic under special QoS requirements, it has to notify the controller about those requirements and also receive the response to know whether some special treatment is granted or not. Fig. 2 illustrates steps of this interaction. For this purpose, in API function calls, before a socket to the destination is created, there is a negotiation phase with the controller in which a request packet containing application and destination identifiers, socket type and intended QoS requirements in the body is prepared. Here is what happens next:

- A TCP socket to the controller is created with packet DSCP bits set to `110110`, a value not in commonly used values in [5] for consistency. We also pre-install a rule after discovery of each switch to match on these specific DSCP bits destined to the controller and forward those packets for further processing (step 1).
- When the controller module responsible for this API receives this type of packet, it prepares a response according to what we described earlier, installs a rule in switches on path to the requesting application and forwards the response (step 2).
- The requesting application finally acknowledges the response or revises the request based on specified limits announced by the controller and acknowledges revision request (step 3).
- The controller installs necessary rules in switches along the path with feasible resources (step 4).
- Based on rules set in step 2, the application-side API module which is waiting for resource allocation completion signal from the controller is notified about this (step 5).
- Finally the application-side module creates its socket to the destination as before and main data flow starts by the application (step 6).

C. Monitoring Network Resources

Generally, there are two approaches to network monitoring and measurement. In active monitoring, some measurement traffic is injected into the network. For example a probe sent to the network will trigger a response and one could analyze the response to infer properties of the network. In passive monitoring, the traffic that naturally exists in the network is observed.

For OpenFlow networks [13], active monitoring translates to relying on `OFPT_MULTIPART_REQUEST` and `OFPT_MULTIPART_REPLY` messages exchanged between switches and the controller. While polling switches for statistics at arbitrarily high rates may provide accurate view of the network, this could incur significant overhead and impose scalability problems in large, heavily loaded networks. On the other hand, passive monitoring in OpenFlow networks means listening to asynchronous messages `OFPT_PACKET_IN`, `OFPT_FLOW_REMOVED` and `OFPT_PORT_STATUS` to infer some properties like link utilization and topology updates. Although passive method has less overhead, but it sacrifices accuracy and also not all QoS metrics can be inferred in this manner.

The monitoring methods plugged into this framework should provide a real-time, accurate snapshot of network so that the framework could properly decide how to respond to applications requests. We utilize adaptive, active statistics collection method similar to the algorithm in [14] which trades between accuracy and overhead and apply it to ports along simple network paths from source to destination. It adapts polling frequency to activity of the ports or amount of traffic that flows contribute to them and calculates link utilization for bandwidth guarantees and traffic shaping and also number of

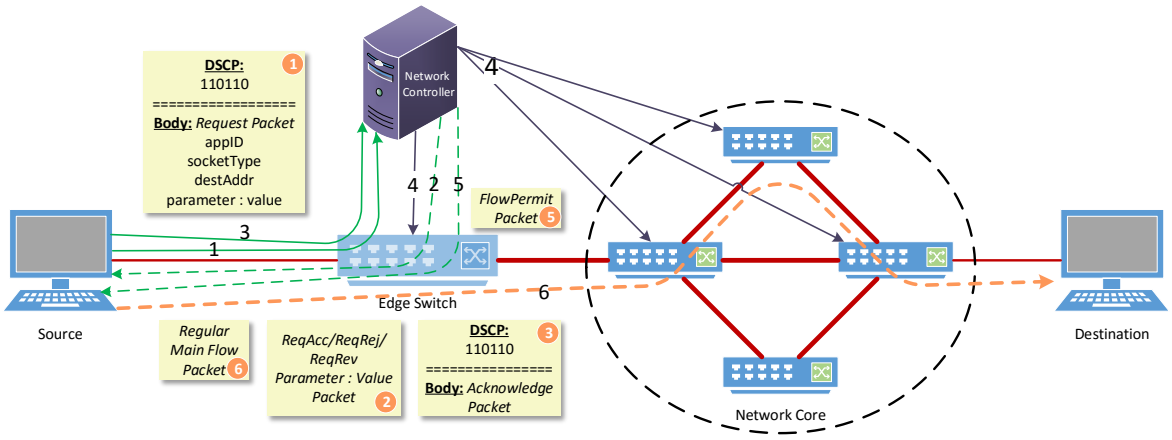

Fig. 2. Steps of application-controller interaction until traffic flow

dropped packets and errors occurred. Regarding path delays, it is not feasible to apriori maintain all-pairs delays. This is because total delay is composed of several components such as processing, queueing and transmission delays which are added up together and cannot be independently captured and saved before a probe destined for that specific path is sent by `OFPT_PACKET_OUT` and `OFPT_FLOW_MOD` messages and also due to possible large number of simple paths between source and destination and limited size of flow tables, such a solution would not be practical. In this case we suggest to set the rules for navigating the probe towards the egress switch and send it upon the receipt of this kind of request.

IV. Prototype Implementation

We implemented bandwidth reservation and traffic shaping scenarios as instances of our proposed framework usage over Ryu which is a widely used component-based SDN controller [15]. Ryu is a modular controller. Each module registers itself to listen for specific events (i.e., messages from other modules) and upon receipt of events, puts them in a FIFO queue. It also has a thread for event processing which keeps dequeuing from this receive queue and calling appropriate event handler. We have a module that tracks topology changes based on LLDP packets plus a monitoring module that maintains links utilization and ports status. Another module processes application requests for validity based on network snapshot information and administration policies and evaluates whether the request complies with its criteria in Table I. For example, regarding minimum bandwidth reservation scenario it checks if the requested amount is less than the maximum bandwidth considered for that application and there is a path which can accommodate the flow. Finally a module implements the behavior in data plane.

We used OpenFlow version 1.3 [13] for communication between the control plane and Open vSwitch [16] data plane in Mininet [17] test and prototyping environment. Mininet uses lightweight virtualization meaning a host is a shell process moved into its own network namespace. Mininet hosts have their own virtual Ethernet interface(s) [18]. For generating test traffic we used iPerf [19] and for graph algorithms we used NetworkX package [20]. All scripts were written in Python.

To implement bandwidth reservation and traffic shaping scenarios we used OVSDB [21] management interface as a complementary protocol to OpenFlow for configuration of switches, ports, queues and attachment of QoS policies to queues. We chose hierarchical token bucket queuing discipline and set configurations in JSON format.

V. Evaluation

We simultaneously evaluated the framework operation under the two implemented scenarios by requesting same upper and lower bandwidth limit through our API functions. A tree topology with 7 switches (`s1-s7`) and 8 hosts (`h1-h8`) was created with Mininet. All link capacities were set to 20Mbps. Two UDP flows `f1` originating from `h2`, destined to `h7` and `f2` originating from `h1`, destined to `h8` were created with iPerf. Flow `f1` was configured as a link throttling flow while `f2` was intended to receive 10Mbps bandwidth. Test environment is depicted in Fig. 3.

The bandwidth perceived by `f2` was measured in two circumstances, with and without calling API functions, to study framework operation and API effectiveness. We configured iPerf to report the bandwidth for each flow in 0.5 second intervals for a 10 second time period. For evaluation purpose, upper and lower bounds bandwidth policies for the application generating `f2` were set in compliance with its request. As shown in Fig. 4, the average `f2` rate is 9.72Mbps with minimum variance when using our framework compared to the average of 5.2Mbps without using the framework. When the connection is made in traditional way, `f1` overwhelms `f2` by competing over shared links. But calling function `reserveMinBW()` before flow `f2` starts, makes the controller set some queues along the path to the destination. For a graph $G = (V, E)$, a simple path is found in $O(V + E)$ using

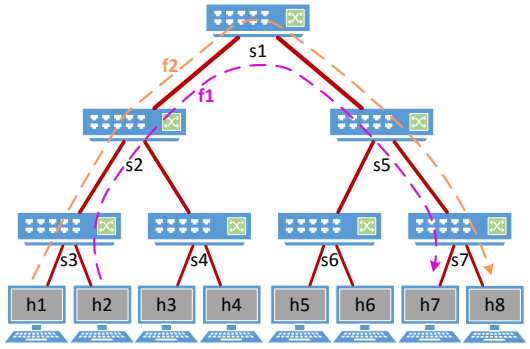

Fig. 3. Topology used for evaluating effectiveness of bandwidth reservation and traffic shaping scenarios. All links have 20Mbps capacity and intended minimum and maximum rate for flow f2 is 10Mbps.

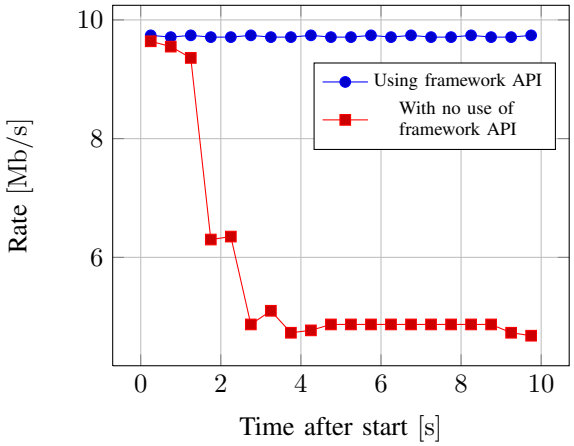

Fig. 4. Bandwidth perceived by UDP flow f2

a modified DFS algorithm but since the number of simple paths might be $O(n!)$ for a graph of degree n. Furthermore, a multi-commodity flow problem arises if requests are not handled readily. Link (u, v) has capacity $c(u, v)$. There are k commodities $K_1, K_2, ..., K_k$ defined by $K_i = (s_i, t_i, d_i)$ where s_i and t_i are source and sick of commodity i and d_i is the bandwidth demand. The flow of commodity i along link (u, v) is $f_i(u, v)$ and the problem is finding an assignment of flows which satisfies the following constraints:

$$\sum_{i=1}^{k} f_i(u,v) \leq c(u,v) \tag{1}$$

$$\sum_{w \in V} f_i(u,w) = 0 \text{ when } u \neq s_i, t_i \tag{2}$$

$$\forall v, u : f_i(v,u) = -f_i(u,v) \tag{3}$$

$$\sum_{w \in V} f_i(s_i,w) = \sum_{w \in V} f_i(w,t_i) = d_i \tag{4}$$

which is NP-Complete while we are facing a time-critical decision. We thus confine to a first-fit heuristic for flow

TABLE II
CONTROLLER PERFORMANCE AND PROCESSING TIME FOR BANDWIDTH RESERVATION SCENARIO UNDER DIFFERENT NETWORK SIZES

Number of Pods (k)	**Number of Edge Switches** $(k^2/2)$	**Average Path Computation Time** $[s]$	**Average Number of New Flows Handled** $[1/s]$
4	8	4.63466085052e−6	2769.29
8	32	2.29313795225e−5	2726.00
12	72	5.93899789978e−5	2588.71
16	128	0.000132252089811	2374.65
20	200	0.000241951589853	2256.31
24	288	0.000388239264119	1967.76
28	392	0.000587952680526	1821.77
32	512	0.000849910279676	1729.95
36	648	0.00125786274914	1302.46
40	800	0.00170276373988	1135.03
44	968	0.00233802724793	686.76

placement. Table II indicates path computation time for flows generated in a k-ary fat-tree network topology introduced in [22]. In this topology which is an instance of Clos topology, there are k pods, each containing two layers of $k/2$ switches with k ports. So there are $k^2/2$ edge switches that we consider triggering points of `OFPT_PACKET_IN` messages. There are $(k/2)^2$ paths between any two hosts on different pods and for evaluation we assume inter-pod traffic, therefore our first-fit heuristic finds paths for bandwidth reservation scenario in $O(k^3)$. We simulate a performance test for the controller by sending a `OFPT_PACKET_IN` from a switch to the controller and waiting for a matching `OFPT_FLOW_MOD` and counting how many times this happens in a second. In each round, we perform this for each one of $k^2/2$ edge switches. We do the test for 15 rounds and average over total number of responses received and path computation times. As evident in Table II, when the number of pods (same as the number of switch ports) increases, the number of switches increases even more and this leads to a larger graph to search and more path computation time. Meanwhile, since a greater portion of controller's time is dedicated to graph computations, each event handling would take more time and less new flows could be serviced in each time period. But even for a network with $k = 40$ which supports $k^3/4 = 16000$ hosts, this time is less than 2 milliseconds.

In order to get a better sense of the overhead our graph computations impose on the controller, we also compared the average number of new flows handled by the controller in two situations, when the controller is barely running a packet_in handler to respond with a simple `OFPT_FLOW_MOD` with no hesitation and when an algorithm of $O(k^3)$ is run by the handler before responding. Fig. 5 demonstrates the results for different network sizes. As depicted in this figure, the controller remains performant even in highest traffic demands. Note that we have considered the worst-case scenario for overhead analysis because the number of times a graph algorithm is required to run, depends on the number of accepted

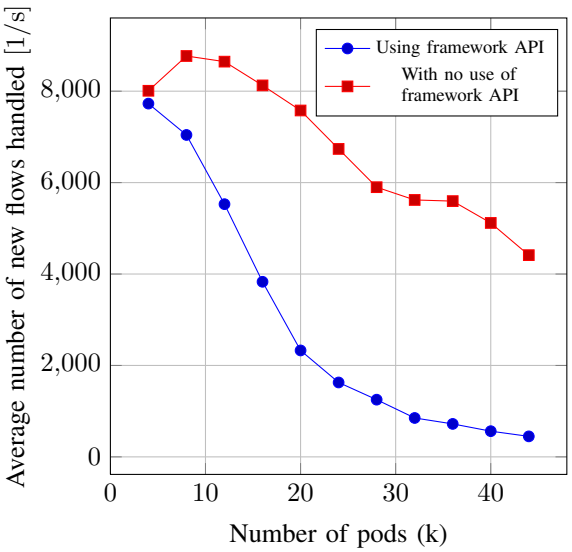

Fig. 5. Controller's ability to handle new flows under different network sizes. Worst-case condition is assumed for API usage.

requests from applications and we know that normally, in a MapReduce [23] shuffle for example, there is a processing time between requests. So, not all packet_in events require graph computations while this is the case with our performance tests. The same reasoning applies to graph search operation where many graph searches terminate sooner than we assumed. Anyway, as mentioned earlier, there is room for improvement of monitoring and controlling algorithms used in this framework when specific knowledge of traffic and communication patterns is present.

VI. Conclusion

This paper describes a framework for delegating traffic management tasks such as QoS enforcement to end-host applications in a software-defined environment. This is a move towards more practical application-aware networking enabled by centralized control in SDNs. With this framework, network administrators could specify policies for authorized applications and network application developers could use an API to inform the controller of their specific requirement for a flow. Several monitoring and control mechanisms can be incorporated into the framework. Evaluation of our implemented prototype for minimum bandwidth reservation and traffic shaping functions shows effective operation of this framework allowing integration of further traffic management tasks and network monitoring approaches.

References

[1] N. Feamster, J. Rexford, and E. Zegura, "The road to sdn," *Queue*, vol. 11, no. 12, pp. 20:20–20:40, Dec. 2013.

[2] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: Enabling innovation in campus networks," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, Mar. 2008.

[3] D. Kreutz, F. Ramos, P. Esteves Verissimo, C. Esteve Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, Jan. 2015.

[4] I. F. Akyildiz, A. Lee, P. Wang, M. Luo, and W. Chou, "A roadmap for traffic engineering in sdn-openflow networks," *Comput. Netw.*, vol. 71, pp. 1–30, Oct. 2014.

[5] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An Architecture for Differentiated Service," RFC 2475, Internet Engineering Task Force, Dec. 1998, updated by RFC 3260.

[6] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: an Overview," RFC 1633, Internet Engineering Task Force, Jun. 1994.

[7] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, "Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification," RFC 2205, Internet Engineering Task Force, Sep. 1997, updated by RFCs 2750, 3936, 4495, 5946.

[8] M. Al-Fares, S. Radhakrishnan, B. Raghavan, N. Huang, and A. Vahdat, "Hedera: Dynamic flow scheduling for data center networks." in *NSDI*. USENIX Association, 2010, pp. 281–296.

[9] A. Curtis, W. Kim, and P. Yalagandula, "Mahout: Low-overhead datacenter traffic management using end-host-based elephant detection," in *INFOCOM, 2011 Proceedings IEEE*, Apr. 2011, pp. 1629–1637.

[10] T. Benson, A. Anand, A. Akella, and M. Zhang, "Microte: Fine grained traffic engineering for data centers," in *Proceedings of the Seventh COnference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT '11. New York, NY, USA: ACM, 2011, pp. 8:1–8:12.

[11] T. Benson, A. Akella, and D. A. Maltz, "Network traffic characteristics of data centers in the wild," in *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, ser. IMC '10. New York, NY, USA: ACM, 2010, pp. 267–280.

[12] W. Kim, P. Sharma, J. Lee, S. Banerjee, J. Tourrilhes, S.-J. Lee, and P. Yalagandula, "Automated and scalable qos control for network convergence," in *Proceedings of the 2010 Internet Network Management Conference on Research on Enterprise Networking*, ser. INM/WREN'10. Berkeley, CA, USA: USENIX Association, 2010.

[13] Open Networking Foundation, "OpenFlow Switch Specification," Mar. 2015, Version 1.3.5 (Protocol version 0x04). [Online]. Available: http://ow.ly/uGXi304Xbee

[14] S. Chowdhury, M. Bari, R. Ahmed, and R. Boutaba, "Payless: A low cost network monitoring framework for software defined networks," in *Network Operations and Management Symposium (NOMS), 2014 IEEE*, May 2014, pp. 1–9.

[15] Ryu SDN Framework. [Online]. Available: https://osrg.github.io/ryu/

[16] Open vSwitch. [Online]. Available: http://openvswitch.org/

[17] Mininet. [Online]. Available: http://mininet.org/

[18] B. Lantz, B. Heller, and N. McKeown, "A network in a laptop: Rapid prototyping for software-defined networks," in *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, ser. Hotnets-IX. New York, NY, USA: ACM, 2010, pp. 19:1–19:6.

[19] iPerf. [Online]. Available: https://iperf.fr/

[20] NeworkX. [Online]. Available: https://networkx.github.io/

[21] B. Pfaff and B. Davie, "The Open vSwitch Database Management Protocol," RFC 7047, Internet Engineering Task Force, Dec. 2013.

[22] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," in *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication*, ser. SIGCOMM '08. New York, NY, USA: ACM, 2008, pp. 63–74.

[23] J. Dean and S. Ghemawat, "Mapreduce: Simplified data processing on large clusters," in *Proceedings of the 6th Conference on Symposium on Operarting Systems Design & Implementation - Volume 6*, ser. OSDI'04. Berkeley, CA, USA: USENIX Association, 2004.